

## EXHIBIT D

February 14th call

Simona Cohen on 24/02/2000 at 11:01  
Category: IGS Meeting Minutes

Attendees:

Jeffrey R Baker/Watson/IBM@IBMUS  
Gary Quesenberry/Raleigh/IBM@ibmus  
Haim Nelken/IBMHAIFA@IBMHAIFA  
Simona Cohen/IBMHAIFA@IBMHAIFA

Minutes: (summarized by Jeff)

We reviewed and discussed the need for the UCR project plan to interlock with SRM. Specifically:

- the UCR pilot exit criteria needs to be established that SRM will use as the basis to move to deploy UCR.
- assuming the UCR pilot meets the exit criteria expected deployment plans are needed
- assuming a successful pilot and deployment, the savings expected, and by whom, during each phase is needed.
- The pilot will probably be with Watson
- Gary will determine items 1, 2, and 3 in the next two weeks. We will meet on a call the middle of next week.

---

February 23rd call

Simona Cohen on 24/02/2000 at 10:57  
Category: IGS Meeting Minutes

Attendees:

Jeffrey R Baker/Watson/IBM@IBMUS  
Gary Quesenberry/Raleigh/IBM@ibmus  
Simona Cohen/IBMHAIFA@IBMHAIFA

Minutes:

We discussed the ROI for UCR, the exit criteria, and how the UCR plan is interlocked with SRM plan. Gary will write all this.

We mentioned the following dates:

- March 15th - Getting access to SRM and the new EPP database.

- March 15th - Getting the revised schemas of SRM and EPP.

- April 30th - Having real data for the pilot selected customer. The data should include EPP probes that operate on servers which are monitored by SRM.

- Mid December - put UCR in SRM test system

- Mid January 2001 - put UCR in SRM production system

The pilot will probably be on Watson notes servers.

In the future, UCR may be appropriate for the NCO project. NCO is migrating from NotesView to Tivoli and they may use EPP as an additional data feed. NCO contact person is Johnny Johns.

Simona will update the UCR plan as follow:

- change dates assuming the project start in a week
- indicate who in Services approved each prerequisite.
- Indicate when review and feedback for deliverables is expected from Services.

Gary said that Toronto is no longer involved in EPP/SRM unification.

Chris is now responsible on integrating the EPP database into the SRM

database. We need to coordinate with Chris whether they rely on UCR only to do the reporting on the unified data.  
Gary said that SRM and EPP may become strategic services in IGS, which will benefit us all.

---

February 28th call

Simona Cohen on 28/02/2000 at 23:28  
Category: IGS Meeting Minutes

Attendees:

Jeffrey R Baker/Watson/IBM@IBMUS  
Gary Quesenberry/Raleigh/IBM@ibmus  
Simona Cohen/IBMHAIFA@IBMHAIFA

Minutes:

The EPP tables are already defined in the SRM database. Now, EPP team is recompiling the load code and then they will start feeding data into the tables.

Johnny Johns was interested about using UCR, so NCO may be our customer if the pilot is successful. NCO includes about 500-600 notes servers, out of them 20-30 are Watson notes servers. The pilot will include the Watson notes servers.

The following issues need to be resolved before the project is approved:

- SRM needs to commit to cost avoidance and savings.

- Customer expected cost savings needed at pilot completion (pursuing WW NCO account).

- Several dependencies interlocked, but all must be.

They will be further discussed in a call on Friday.

---

SRM & UCR Agreement

Simona Cohen on 05/03/2000 at 13:04  
Category: ISO

To: Jeffrey R Baker/Watson/IBM@IBMUS  
cc: Simona Cohen/Haifa/IBM@IBMIL, Johnny Jones/Raleigh/IBM, Kim Hansma/Lexington/IBM@IBMUS  
From: Gary Quesenberry/Raleigh/IBM@ibmus  
Subject: SRM & UCR Agreement

Jeff,

Based on our meetings UCR project meetings to-date, we agreed that the following issues needed to be resolved before the project was formally approved. I have documented our commitments and the current status with each issue.

SRM needs to commit to cost avoidance and savings. The SRM development team commits up to 0.2 FTE savings (based on a small core development base to begin with) upon successful integration with SRM; we also commit up to 1 PM development cost avoidance savings for each new data metric that SRM will add in the future, such as Middleware metrics, etc... The SRM team will also contribute the necessary project management, architecture, and SRM (GUI, database) development hours in order to integrate the UCR solution with SRM 4.6 scheduled for delivery January/2001.

Customer expected cost savings needed at pilot completion (pursuing WW NCO account). Any customer savings, customer requirements, and feedback will be coordinated through Johnny Jones' NCO Project Office team. The current worldwide NCO install base which incorporates both SRM and EPP deliverables today, will be a natural fit for demonstrating UCR in a production environment. Follow-on demonstrations and team meetings will refine the customer requirements and expectations.

Technical pre-reqs The following technical pre-reqs documented in the UCR Plan will be executed by the SRM and IGS teams:

- Assistance in establishing connection to SRM and EPP databases.
- Assistance in understanding the new SRM and EPP schemas.
- Assistance in defining requirements/specifications.
- Assistance in design review (optional).
- Selecting the pilot environment and customer.
- Assistance in deploying the pilot in Raleigh.

I hope this is what you are looking for as to commitment from our end. Let me know if anything else needs to be added or refined. Thanks.  
....Gary

Gary Quesenberry  
Certified I/T Specialist  
DSM Server Resource Management  
IBM Global Services, SDC - South  
919-254-1217 or T/L 444-1217

---

## XSLT

Simona Cohen on 07/03/2000 at 13:35  
Category: Resources

The standard - <http://www.w3.org/TR/xslt>

A good tutorial: <http://metalab.unc.edu/xml/books/bible/updates/14.html>

Apache open source XSLT engine: <http://xml.apache.org/xalan/overview.html>

Many examples: <http://zvon.vscht.cz/HTMLonly/XSLTutorial/Books/Book1/index.html> or

<http://zvon.vscht.cz/HTMLonly/XSLTutorial/Books/Book1/bookInOne.html> (no frames)

An example how to use XSLT to transform from XMI to HTML -  
[http://www.objectsbydesign.com/projects/xmi\\_to\\_html.html](http://www.objectsbydesign.com/projects/xmi_to_html.html)

---

## Feedback on Demo from Johnny Jones (NCO)

Simona Cohen on 08/03/2000 at 15:39  
Category: Demo

Here is a summary of the feedback. My answers are in blue.

There is a different probe ID for each day. How can we see the same probe running over days?

We can categorize the probes by a metric that will put all the same probes for different days in one layer (e.g. Target ID).

We also need to think on a better way.

There are occasionally null pointer exceptions when drawing the white

flyover window.

We need to decrease the amount of memory used.

How is the NotesView data (e.g. number of concurrent users) going to be displayed?

This is the application data, so it should be probably displayed on the flyover of the probes.

These metrics will also be added to the "Color Probes By" listbox, so we can change the color of probes according to the new metrics.

In the back-end we create all the unified data. What if someone wants to view just part of the report.

In case of an on-line (dynamic) report where the unified data is created at the time the report is requested, it is a waste to create all the unified data.

This should be considered in the performance improvements in 4Q.

---

#### UCR vs. Garlic

Simona Cohen on 13/03/2000 at 13:22

Category: Back End System, Similar Products

Garlic initial intention was to allow search/query in multi-media objects, less on the correlation and consolidation.

I see the following differences between UCR and Garlic:

standards - UCR uses DTD/XML to define the unified schema while Garlic uses GDL (Garlic Data Language). UCR uses XSLT (probably) to define the mappings while Garlic uses its own C++ API. UCR will use XML query language to pose queries on the unified data while garlic uses GQL (Garlic Query Language).

multi-platform - Garlic was written in C++ for RS/6000. UCR is written in Java.

technology - Garlic is focused on on-line unification only and optimization of GQL queries. In many cases, we would like to have both a data warehouse of the unified data and on-line query on the diverse data sources. Historical data is one such case where we would like to access the data warehouse for "last week data" and on-line query for "today data". The ASDP impact analysis also assumes a data warehouse for the unified data. UCR provides both options.

mapping - UCR back-end provides dynamic creation of a chain of mapping/conversion functions and is not limited to one stage conversion. This can be an advantage when there is a need to use existing conversions or when need to add functionality (such as conversion to voice, other languages, additional formats, exports etc.).

---

#### March 14th meeting

Simona Cohen on 14/03/2000 at 14:52

Category: Meeting Minutes

#### Attendees:

Leonid Dubinsky/Haifa/IBM@IBMHAIFA

Haim Nelken//Haifa/IBM@IBMHAIFA

Hilit Grosberg/IBMHAIFA@IBMHAIFA

Simona Cohen/IBMHAIFA@IBMHAIFA

Minutes:

We'll have weekly meetings on Tuesdays at 14:00 in room 504

We are invited to TLC2000 that will be on April 2-5 in Austin, Texas to demonstrate UCR.

Simona had a call with Martin in which we agreed that

(a) Martin will write a document to describe his code.

(b) Martin is not developing CREDiT anymore, so there is no must to coordinate code with him.

(c) Bernice is working on a new "bigger and better toolkit" based on Diamond that will integrate with CREDiT but the schedule of this toolkit is not clear. Prina will talk with Bernice about co-operation on that toolkit.

Hilit is working on doing transformations using XSLT and will write a document to define the types of transformations that we want to have.

Leonid will try the new applet download time and work within IE. He is also working on technology transfer.

Simona will work on DOU.

Simona will work on DB2 Connect problems.

We all work on specifications document.

---

CRM Data Engine

Simona Cohen on 19/03/2000 at 15:46

Category: Back End System

I've looked into the CRM Data Engine Specification as taken from the Quick Place database, in order to check whether we can take advantage of it.

---

March 21st meeting

Simona Cohen on 21/03/2000 at 14:50

Category: Meeting Minutes

Attendees:

Leonid Dubinsky/Haifa/IBM@IBMHAIFA

Hilit Grosberg/IBMHAIFA@IBMHAIFA

Simona Cohen/IBMHAIFA@IBMHAIFA

Minutes:

DB2 Connect problem was solved by adding MTU in the registry. So, now we can access all data of SRM. EPP data is not available yet.

We do not have an AIX machine yet.

Martin sent us initial documentation on CREDiT and Leonid will try to drill down to subprobes as result of user event instead of loading all subprobes at the beginning.

Hilit will continue working on using Xalan (XSLT processor) to do transformations from data sources to unified schema.

Simona worked on doing functions on selected node-sets with XSLT. For example, write an `Server_performance_info` element for the average `mem_percentage_used` per month.

DOU draft is almost ready.

In the coming week, we all work on specifications draft: Leonid -front end, Hilit-administrator, Simona-overall and back end.

---

#### DOU 2000 Review

Simona Cohen on 27/03/2000 at 11:58  
Category: ISO

##### Attendees:

Oded Cohn/IBMHAIFA@IBMHAIFA  
Haim Nelken/IBMHAIFA@IBMHAIFA  
Simona Cohen/IBMHAIFA@IBMHAIFA

##### Minutes:

Rephrase the first item in 7.1 to a positive statement, and emphasise that not all the potential of the project can be seen in the pilot.  
Split the second item in 7.1 to (1) cooperation with customer and (2) novice GUI.  
Remove 7.3  
We went over the contract review checklist - Q52 and all the items were OK.

---

#### March 28th meeting

Simona Cohen on 28/03/2000 at 13:48  
Category: Meeting Minutes

##### Attendees:

Leonid Dubinsky/Haifa/IBM@IBMHAIFA  
Hilit Grosberg/IBMHAIFA@IBMHAIFA  
Simona Cohen/IBMHAIFA@IBMHAIFA

##### Minutes:

Leonid is studying CREDiT code and changed that subprobes will be downloaded as a drilldown event instead of at the beginning. This was done for application and now this is done for applet.  
Hilit made a version of our XSL example with the open source Xalan and Xerces. Now she will try adding extensions to our example i.e. an XML element/attribute in the input is processed by a Java conversion function to produce the output element/attribute.  
We are all working on the spec, and we'll schedule review for next Tuesday.  
The DOU was reviewed yesterday and sent to Jeff Baker.  
Martin told us that he is leaving IBM this week, but he will give us his email in his new job.  
Vladimir will represent us in TLC2000.

---

#### Specifications discussion - April 4th

Simona Cohen on 05/04/2000 at 13:36  
Category: ISO, Meeting Minutes, UCR Specifications

##### Attendees:

Hilit Grosberg/Haifa/Contr/IBM@IBMIL  
Haim Nelken/Haifa/IBM@IBMIL  
Ilan Paleiov/Haifa/IBM@IBMIL  
Noga Meshulam/Haifa/IBM@IBMIL

Yoram Adler/Haifa/IBM@IBMIL  
Yuri Levin/Haifa/IBM@IBMIL  
Simona Cohen/Haifa/IBM@IBMIL  
Leonid Dubinsky/Haifa/IBM@IBMIL (only front-end)  
Prina Vortman/Haifa/IBM@IBMIL (only back-end)  
Opher Etzion/Haifa/IBM@IBMIL (only back-end)  
David Botzer/Haifa/IBM@IBMIL (only back-end)  
Tali Yatzkar-Haham/Haifa/IBM@IBMIL (only back-end)

Minutes:

In this meeting we discussed the first draft of UCR specifications document, just before sending it to IGS.  
We went through most of the major section - section 4 "Required Functions Specification".

General :

To find the "killer report", we may want to talk with Alex Shmugliakov and ask him what types of reports he needs.  
Ilan and Opher will send us reports from Tivoli, so we get a better feeling of what types of reports are used.  
We should find out how long SRM saves information in the database. Do they save history for a year? Is it longer than that?  
We should find out what is the purpose of the reports. Is it problem determination? Is it capacity planning? Is it something else?

Section 4.1 - UCR Access From SRM Site

The name Correlation for the new button is not good because the meaning is not clear. We need to think on a better name.  
In "Report Type" we should add "All" for a report that includes all probes done on the selected servers.  
The timezone is an important issue and many applications fail on this.  
We should consider handling this, and give a comprehensive solution from the beginning.

Section 4.5 - Back-end Specifications

To make the back-end more generic, we should not limit the unified data to be in XML. Also, we should not limit the mappings language to be XSLT. This can be done by adding more abstract interfaces.

We should look into the multimedia transformation work in which the input parameters for a transformation were:

- domain - e.g. voice, data, image
- type - e.g. mpeg, au
- version - mpeg3, mpeg2
- language - English, Hebrew
- conversion - a dll or a jar

Prina will send us more information on the multimedia transformation work.

We should make our back-end abstract enough to capture our data transformations as well as the multimedia transformations.

For scalability reasons, we should check whether to use EJB. We may check that with Eliezer Dekel and Yoad Gidron.

We should add that the conversion functions can also have the duty of cleaning the data that is in the data sources.

The mapping language should be able to handle a many-to-many mapping. In the specifications for the mapping language, we should add more examples to make it clearer.

Section 4.4 - Administrator Specifications

We should state clearly that not all mappings can be done by visual means only.

When doing a mapping, we may add a link (line). Then when firing the line, a text area comes up where the user can add XSLT code or parameters for the conversion functions.

The XSLT code entered and the conversion functions should be checked for correctness.

We should handle collision between mappings.

We should use colors to ease the creation of mappings.

We may use a CREDiT view to select a view of the data source e.g. the rows from perform.server that belong to a specific fqhn.

Sometimes we have an element that contains the same information like another element and the second is a drill-down of the first. For example the server resources per day and the server resources per hour. In such cases, there is no need that the administrator user will enter all the mappings again for the second object. We should have some automation here.

#### Section 4.4 - Visualization Applet Specifications

The CREDiT view has the following limitations:

We can see only one dimension from each data source at a time. For example, we can not see cpu utilization and memory used together.

It is hard to see small changes because it's a small change in the color.

It's hard to see that every day at a specific hour there is degradation of performance.

A new view should solve the above.

In CREDiT, you cannot see the information from all the fly-overs together. It was commented that there is no need for such a view. We need to check that.

It is important to have a window of all the problems and ease the spotting of problems.

To decrease download time, we may consider using a plug-in.

When there is a lot of data to be put in the fly-over, we may put only a small menu in the fly-over and then give the details per request. This is taken from the Windows 2000 UI.

Check whether we can support IE4.

We should have a button that export the data to an EXECLL or some other relevant application.

The Reviewer Notes can be either as sticky-pad that are stacked to the area of interest, or can be an overall window.

---

April 11th call

Simona Cohen on 13/04/2000 at 08:41

Category: IGS Meeting Minutes, UCR Specifications

#### Attendees:

Jeffrey R Baker/Watson/IBM@IBMUS  
Gary Quesenberry/Raleigh/IBM@IBMUS  
John Froehlich/Lexington/IBM@IBMUS  
Amy Hahn/Lexington/IBM@IBMUS  
Johnny Jones/Raleigh/IBM@ibmus  
Hilit Grosberg/Haifa/Contr/IBM@IBMIL  
Leonid Dubinsky/Haifa/IBM@IBMIL  
Simona Cohen/IBMHAIFA@IBMHAIFA

#### Minutes:



In this meeting we mostly reviewed the UCR Specifications draft 1.0 - (Document link: Database 'Unified Customer Reporting', View 'By Category', Document 'FSD - draft 1.0').

We answered the questions in it as follows:

The purpose of the reports is mainly low level problem determination. Currently we are going to correlate EPP data and system resources data only. We are not going to correlate NotesView data (e.g. number of users) as it about to change. In the future, we may revisit the plan and consider adding Notes application specific data probably taken from Tivoli.

We'll have a talk with Johnny's team to further determine the purpose of the reports and the "killing report".

An NCO report is SRM includes today 70 servers and it can easily grow to 100. This means that the CREDiT applet will consume a lot of client memory which is a problem. Moving to a more servlet-based approach means that the user will have to wait a lot each time he wants to color the probes or servers by another dimension/field.

Haifa will try to solve the problem. Then, it may also be used by SDTC ESMA enterprise data.

In case it is not solved, some ideas were raised:

- Show just the red and yellow servers.
- Give a report for one server at a time but much more detailed.

This is for problem determination.

-- Have something like sub\_client\_id3 that will further divide an account to states/cities.

In the pilot we are not going to take in consideration the shift. The information is summarized per day.

In SRM, the time used is that at the server. If the time in EPP is according to the station timezone, some matching need to be done. We need to know the GMT offset for the station and the server. We need to further investigate this with Andy.

All the EPP-SRM correlation questions will be submitted to Andy.

For now, there is no server type in SRM. SRM will add it in the future when there will be an accurate inventory database.

We don't have to list all CPUs and disks.

The overall status of a server is the worst case for one of its resources. In other words, if a server is yellow in CPU, green in memory and red in disk - it's a red server.

The server "general load" algorithm need to be obtained from Chris.

SRM is adding now separate thresholds for each server.

we need to ask Andy if we miss some of the data in the table of reported fields.

There will be no NLS support in the pilot. We'll support just English.

We discussed some other issues as follows:

EPP data will probably go into the tables this week.

On April 25th we have the Passover holiday in Israel, thus we won't have neither weekly call nor monthly status. I'll send a status note on April 27th.

Simona will update the specifications document (FSD) as soon as we get answers from Andy.

---

April 18th meeting

Simona Cohen on 18/04/2000 at 18:34

Category: Meeting Minutes

Attendees:

Hilit Grosberg/IBMHAIFA@IBMHAIFA  
Simona Cohen/IBMHAIFA@IBMHAIFA

Minutes:

Leonid and Noga are in vacation.

Simona has discussed EJBs with Yoad and Gal and summarized here -  
(Document link: Database 'Unified Customer Reporting', View 'By  
Category', Document 'EJBs - talk with Yoad and Gal'). We decided not to  
use it.

A meeting with NCO people will be on May 3rd. We'll get feedback on the  
demo and FSD.

We got DB2 Connect 6.1 for NT and AIX.

Hilit worked on XSLT - extension that returns a node list, the Redirect  
extension, several XMLs in the input.

We'll have a design meeting on May 11th and need to send document to IGS  
by May 15th.

No meeting next week due to Passover.

---

FSD - release 1.1

Simona Cohen on 30/04/2000 at 15:56  
Category: ISO, UCR Specifications

Here is the UCR specifications release 1.1:  
(See attached file: ucrfsd1.1.zip)

It is also available at <http://9.148.34.57/UCR/docs/ucrfsd.html>

In this version, I've made changes to sections 1.5, 2.4, 4.0, 4.1, 4.2,  
4.3, 13.2 according to your input.

---

May 2nd call

Simona Cohen on 02/05/2000 at 17:03  
Category: IGS Meeting Minutes

Attendees:

Jeffrey R Baker/Watson/IBM@IBMUS  
Gary Quesenberry/Raleigh/IBM@IBMUS  
Chris Molloy/Raleigh/IBM@IBMUS  
John Froehlich/Lexington/IBM@IBMUS  
Johnny Jones/Raleigh/IBM@IBMUS  
Andy Frenkiel/Watson/IBM@IBMUS  
Simona Cohen/IBMHAIFA@IBMHAIFA

Minutes:

FSD release 1.1 was sent and no comments on it.

The perform.r\_day\_cpu table is used for graphs and should not be used by  
UCR.

The disk information should not be taken from perform.r\_day,  
perform.r\_week, perform.r\_month but from perform.r\_disk\_day,  
perform.r\_disk\_week, perform.r\_disk\_month respectively. More disks  
details can be found in perform.r\_pdisk\_day (physical disks) and  
perform.r\_ldisk\_day (logical disks).

There is no need for a general\_load field to specify the server health

in numbers. The status field that specify the server health in colors is enough. The new SRM release that went into production last week lets the customer enter its own thresholds.

Also an EPP color will be added for the response time status.

The EPP data is going to be loaded to the SRM database in a week. If it won't be by next week, we may FTP the exported file from the Poughkeepsie database.

The list of Watson Notes servers as mentioned in the FSD is not correct.

We may find an updated list of the Watson notes servers in the Server pull down menu at

<http://corona.watson.ibm.com/wisqos/notesPilot.htm>

---

#### May 9th call

Simona Cohen on 09/05/2000 at 17:36

Category: IGS Meeting Minutes

#### Attendees:

Jeffrey R Baker/Watson/IBM@IBMUS

Gary Quesenberry/Raleigh/IBM@IBMUS

John Froehlich/Lexington/IBM@IBMUS

Andy Frenkiel/Watson/IBM@IBMUS

Simona Cohen/IBMHAIFA@IBMHAIFA

#### Minutes:

John has now access to the tables in Boulder but it seems that the schema is not quite the same as in SRM database. For example, in the transaction table (the primer table), there is a column defined as float in SRM and char(10) in Boulder. John will work on this with Keith W. Kempke. After this is resolved, the EPP data will be dumped from Boulder to Ralieg.

The current schema still misses: (1) lookup tables for unique probes IDs (2) a new table that maps targetID to fqhn.

Haifa will get access to the Boulder MVS DB2 database. The contact person for this is Keith.

Jeff noted that after the demo to NCO, it seems they are receptive to the solution. Yet, a lot of people participated, but only one or two provided feedback. We didn't really got the answer to what is the "killing report". Only after the NCO team uses UCR, more feedback will come.

Scaleability will be probably one of the issues that will arise as soon as UCR is used by NCO, and we should take it in consideration in advance.

---

#### May 16th call

Simona Cohen on 18/05/2000 at 21:17

Category: IGS Meeting Minutes

#### Attendees:

Jeffrey R Baker/Watson/IBM@IBMUS

Gary Quesenberry/Raleigh/IBM@IBMUS

John Froehlich/Lexington/IBM@IBMUS

Johnny Jones/Raleigh/IBM@IBMUS

Simona Cohen/IBMHAIFA@IBMHAIFA

Minutes:

In order to access the EPP data in Boulder as fast as possible, UCR will get the SRM ID to Boulder, and at the same time will continue asking for an ID for UCR.

SRM and EPP need also work on the new table that correlates fqhn to target\_id, including a process to populate this new table.

SRM 4.4 will be delivered next quarter and include many new features as new agents for OS/390 and AS/400, dynamic reports in the internet site, using Tivoli data instead of NotesView, MVS reports, Linux reports, middleware reports (Oracle statistics, web statistics), use of JChart as the graphics toolkit.

In the next quarter EPP and SRM will also deliver a prototype that includes charts of server performance and response time.

Johnny said that having a report with a big list of servers is valuable, but if this is not possible we can have a model where we first show 50 servers, then bring another 50, and so on. Alternately, we can bring servers according to their types- MTA servers, Mail servers, Database servers, Mail Hub servers.

Haifa is starting now, after the design phase to put some of the applet processing in a servlet, only after this is done, we'll have a better idea on how much data we can put at a time.

UCR will be presented in next Service Delivery/Research in-person council meeting by Gary and Haim. Simona will start the work on the presentation template, and will circulate that.

---

May 23rd call

Simona Cohen on 23/05/2000 at 23:34  
Category: IGS Meeting Minutes

Attendees:

Jeffrey R Baker/Watson/IBM@IBMUS  
Amy Hahn/Lexington/IBM@IBMUS  
John Froehlich/Lexington/IBM@IBMUS  
Andy Frenkiel/Watson/IBM@IBMUS  
Simona Cohen/IBMHAIFA@IBMHAIFA

Minutes:

Haifa was able to connect to the Boulder EPP database and view data for the Watson notes servers using SRM ID. We sent to Keith the Simona's employee number in order to get an ID for UCR.

Scott from SRM will send us next week a package that will allow us to build an SRM mini-site here in Haifa for testing purposes. We'll integrate the UCR code into this.

The UCR design draft document was delivered last week. Comments are welcomed!

The people in Boulder need to do some last changes in the EPP schema..

After this is done, we'll get the new schema and the data will be propagated from Boulder to the SRM DB. Also, in a month or so, foreign keys will be added to the tables in Boulder.

A new table was defined for the SRM-EPP correlation to include fqhn, target\_id, dom\_id, mapping\_start\_day, mapping\_end\_day. Andy will send Haifa the definition of this table as well as other changes to the schema when occur in the future.

---

May 30th call

Simona Cohen on 31/05/2000 at 02:17  
Category: IGS Meeting Minutes

Attendees:

Jeffrey R Baker/Watson/IBM@IBMUS  
Gary Quesenberry/Raleigh/IBM@IBMUS  
Amy Hahn/Lexington/IBM@IBMUS  
John Froehlich/Lexington/IBM@IBMUS  
Johnny Jones/Raleigh/IBM@IBMUS  
Haim Nelken/IBMHAIFA@IBMHAIFA  
Simona Cohen/IBMHAIFA@IBMHAIFA

Minutes:

The presentation for the Service Delivery/Research in-person council meeting was reviewed by Haim and Gary and sent to Jeff. UCR ID for the Boulder EPP database doesn't work. We use SRM ID. The data in the Boulder database for the pilot servers is incomplete.. A note with the details will be sent to Andy, and if necessary schedule a call.

The delay in accessing the data and the fact that it is currently partial, will effect the UCR plan. It will delay interim milestones in about a month. We'll also add a pre-req of building an SRM test environment in Haifa. The revised plan will be co-ordinated between Gary and Simona after we get more information on the missing data.

---

June 6th call

Simona Cohen on 07/06/2000 at 21:41  
Category: IGS Meeting Minutes

Attendees:

Amy Hahn/Lexington/IBM@IBMUS  
John Froehlich/Lexington/IBM@IBMUS  
Johnny Jones/Raleigh/IBM@IBMUS  
Keith Kempke/Boulder/IBM@IBMUS  
Simona Cohen/IBMHAIFA@IBMHAIFA

Minutes:

UCR does have now its own ID to Boulder.  
John will follow up with Scott about the testing environment in Haifa.

The following minutes are about the data in the Boulder DB as we learned from Keith

The only valid target\_id for the Watson notes servers are those in the 400

All transactions with test\_id=0 are exceptions and will be removed  
d01ml017 is not probed lately, and we should check that with Drew from Watson

The summary tables will contain Watson servers soon

There is no intention to put data in EPP\_TEST\_OBJ\_TIME and EPP\_TEST\_OBJECTIVE soon. These tables are for future use.

The final schema will be distributed on Friday, and after that the data can be exported from the Boulder DB to the Raleigh DB

---

June 13th call

Simona Cohen on 14/06/2000 at 00:07  
Category: IGS Meeting Minutes

Attendees:

Gary Quesenberry/Raleigh/IBM@IBMUS  
Amy Hahn/Lexington/IBM@IBMUS  
John Froehlich/Lexington/IBM@IBMUS  
Johnny Jones/Raleigh/IBM@IBMUS  
Keith Kempke/Boulder/IBM@IBMUS  
Andy Frenkiel/Watson/IBM@IBMUS  
Simona Cohen/IBMHAIFA@IBMHAIFA

Minutes:

We learned from Drew notes that 7 servers of the pilot are going to be removed by next week. So, we are left with 3 servers, out of which one server (d01ml244) is not monitored by SRM.

Gary will follow up with monitoring d01ml244.pok.ibm.com

Andy will follow up with Drew in defining new probe - NotesReplic on the same 3 servers, so UCR has more data to report on.

At the same time, we'll check the possibility to use another set of servers in the pilot. These servers include email probes that have about 6-10 sub-transactions according to the intermediate mail hubs. Andy said that these probes are complicated. Andy will send us information about these probes and UCR will learn the new data and its impact.

Note that d01ml243.ibm.com from SRM is the same server as d01ml243 in EPP. The pok is missing in the FQHN because of a typo.

---

June 20th call

Simona Cohen on 21/06/2000 at 01:00  
Category: IGS Meeting Minutes

Attendees:

Jeffrey R Baker/Watson/IBM@IBMUS  
Gary Quesenberry/Raleigh/IBM@IBMUS  
Amy Hahn/Lexington/IBM@IBMUS  
Simona Cohen/IBMHAIFA@IBMHAIFA

Minutes:

We agreed that for the next delivery (July 31st) we'll have:

Correlation of real data for the servers that are left from the list defined in the UCR spec - d01ml233, d01ml243, d01ml244.

Correlation of dummy data for as much as possible servers and probes.

This is in order to test the scalability of the applet. It doesn't test however scalability of the back-end.

The EPP data is loaded in Ralieg. It's not clear yet how often this data is updated, so we'll continue using the Boulder database. The correlation table is in Ralieg.

Fred Newel was in vacation and will start monitoring the 3rd server (d01ml244) for SRM this week.

Boulder has many new tables for EPP. We'll ask Keith how to get documentation on those new tables.

Scott will sent us the package to build a testing environment in Haifa tomorrow.